

DEPARTMENT: VISUALIZATION VIEWPOINTS

Using Artificial Intelligence to Visualize the Impacts of Climate Change

Alexandra Luccioni, *Université de Montréal, Mila, Montreal H2S 3H1, Canada*

Victor Schmidt, *Université de Montréal, Mila, Montreal H2S 3H1, Canada*

Vahe Vardanyan, *Ubisoft, Montreal H2T 1S6, Canada*

Yoshua Bengio, *Université de Montréal, Mila, Montreal H2S 3H1, Canada*

Public awareness and concern about climate change often do not match the magnitude of its threat to humans and our environment. One reason for this disagreement is that it is difficult to mentally simulate the effects of a process as complex as climate change and to have a concrete representation of the impact that our individual actions will have on our own future, especially if the consequences are long term and abstract. To overcome these challenges, we propose to use cutting-edge artificial intelligence (AI) approaches to develop an interactive personalized visualization tool, the AI climate impact visualizer. It will allow a user to enter an address—be it their house, their school, or their workplace—and it will provide them with an AI-imagined possible visualization of the future of this location in 2050 following the detrimental effects of climate change such as floods, storms, and wildfires. This image will be accompanied by accessible information regarding the science behind climate change, i.e., why extreme weather events are becoming more frequent and what kinds of changes are happening on a local and global scale.

Historically, climate change has been an issue around which it is hard to mobilize collective action, despite the threat that it presents to our species and life on our planet. Researchers studying this phenomenon have systematically observed that effective communication on the subject of climate change arises from messages that are both emotionally charged and personally relevant, and that *images* in particular are key in increasing the awareness and concern regarding the issue of climate change.²⁷ Despite this assertion, many of the traditional forms of communication that are used by experts to communicate to the public are often based on scientific reports, which can fail to communicate the urgency and importance of this monumental phenomenon to the general public.⁴ The aim of our

project is to use Artificial Intelligence (AI) to contribute to bridging this gap, and to create a tool to raise awareness with regards to the impacts of climate change.

CLIMATE CHANGE COMMUNICATION

Given the findings on the importance of images in climate change communication, recent research has explored translating numerical climate models into representations that are more intuitive and easier to understand, for instance via climate-analog mapping, which involves finding a city whose current climate matches that of another city in 2080, following the impacts of global warming: for example, based on climate model projections, Washington, DC, will have a future climate that is more similar to that of Greenwood, Mississippi, which is 6.4 °F (3.5 °C) warmer and 1.7% drier.⁸ Other projects have endeavored to facilitate more emotional and experiential engagement via a compelling textual

narrative, emphasizing climate change as a present, local, and personal risk.²⁵ Image-centric approaches have also focused on selecting relevant photographs to represent the extent of climate change impacts^{6,23} as well as on using artistic renderings of possible future landscapes⁹ and even immersive video games² and virtual reality experiences.¹

The goal of all of these approaches is to communicate effectively regarding this all-encompassing global phenomenon that will inevitably impact every single living being on the planet. However, therein lies the issue—the impacts of climate change cannot be predicted perfectly, and this is reflected in uncertainty in the climate models and in their probabilistic output. Therefore, communications around climate change have continued to emphasize the *possible* changes that *may* happen in the future, based on numbers from scientific studies. However, as long as this message stays abstract, it is difficult for most people to fully incorporate this information and act accordingly on it. This is due to cognitive biases that we have as human beings, such as those that predispose us to maintain the status quo and to avoid expending resources on threats that are outside our realm of sensory experience, such as those that are far away in time and space.¹⁵ While these biases were an advantage in environments where threats to our survival were concrete, nearby, and impending, such as a woolly mammoth threatening to attack, they are much less useful in fighting current threats to our existence, such as invisible pollution, slowly rising seas, an increase in zoonotic diseases, or added degrees of global warming temperatures. We therefore see a need to use modern, 21st century tools such as AI to overcome these biases and spur action on climate change.

IMAGES AS COMMUNICATION VECTORS

Research in behavioral science has systematically shown that people's behaviors and attitudes can be affected by interactions with images that visually depict the future consequences of actions taken today. For instance, researchers have found that young people tend to save significantly more money when presented with a decision calculator that projects their pension amount, accompanied with an age-progressed photograph of themselves at retirement age.¹² Likewise, middle-aged and older employees commit to putting aside more money for their retirement when the total amount is made more vivid and concrete, as a monthly amount to live on, with concrete examples of the objects that they can and cannot purchase.¹⁰



FIGURE 1. Example of an urban flooding scene generated by our generative model.

The above studies show that when people are given evocative depictions of how actions today affect outcomes tomorrow, they are more likely to change their behavior. In particular, they often behave more like people who have thought about the issues at length. This emotionally grounded “decision aiding” approach maintains credibility by drawing on unbiased projections of the future and leaving freedom of choice to the (now better informed) individual. Conscious of the fine line between educating (i.e., by providing transparent and truthful information) and manipulating (i.e., using advertising or misinformation to influence outcomes), we propose to harness this approach, allowing viewers to time-travel into the future and visualize the severe impacts that climate change is likely to have by using an AI technology.

THE POTENTIAL OF AI

In recent years, AI has made incredible progress in many domains and applications, including speech recognition, machine translation, and computer vision. However, one of the truly impressive applications of AI has been the ability to create incredibly realistic images of people and objects using generative adversarial networks (GANs).¹¹ The most attention-grabbing applications of this generative technology have been the so-called “Deep Fakes,” such as aging and gender transformations and even video editing applications applied to any user-submitted image, with surprisingly realistic results.

Above and beyond these viral examples of GAN applications, there have also been projects that aim to use generative models for a clear social good, for instance to aid people's understanding of and empathy toward situations that affect people and places far



FIGURE 2. Before and after example of the transformation carried out by our GAN approach, based on an image of a suburban house.

away in time and space. For example, MIT's *Deep Empathy project* aims to use AI-generated imagery simulating the impacts of war on major global cities like Boston and Paris in order to help viewers empathize with victims of war in places like Syria and Afghanistan. Nonetheless, research on the psychological and emotional impact of AI-generated imagery is limited, and warrants further empirical studies regarding the ways in which they can impact human perception and decision-making, while addressing the specific moral and technical challenges raised by the creation and usage of AI-generated images.¹⁶

As a matter of fact, the true utility of these generative models arises from their ability to generalize and to produce diverse, realistic outputs. This is especially true in cases where a target output does not exist in real life, making it impossible to simply train using more traditional, "supervised" AI techniques. Generative approaches allow the user to control the generation process, for instance, by defining specific modes, class

labels, or features that the output image should have. These conditional generative models, which can generate an image given some desired characteristics, have notably been used for style transfer, applying artistic styles such as that of van Gogh or Monet to photographs, showing what it could look like if Monet painted the harbor of modern Shanghai or van Gogh the savannah near Nairobi. We harness the same kind of AI generative power in our own project, but instead of applying an artistic style, we illustrate rising waters in landscapes where flooding seems like a far-away impossibility now.

OUR APPROACH

As we stated above, our project aims to raise public awareness and understanding of climate change and to inspire people toward sustainable actions and lifestyle choices. By using AI-generated imagery to visualize the personalized, street-level consequences of climate change, we make the threat of climate change more salient and concrete. The extreme weather event that we chose to focus on first is flooding, since we are based in Quebec, Canada, a region that is increasingly impacted by flooding due to climate change. We believe that using these images will help to overcome the traditional climate change communication and action barriers and raise awareness regarding this important issue. Our hope is that this will enable helping viewers overcome their cognitive biases and to take action on stopping climate change as individuals and as members of the global community.

METHODOLOGY

To generate the images of the locations under the effect of climate-change-induced flooding, we use the same kind of generative approach described in the previous section, but we have created a custom GAN to better suit our purposes. In essence, GANs are AI models that are made of two subnetworks: a generator and a discriminator. These networks compete in a zero-sum game: the goal of the generator is to fool the discriminator by creating images that are as realistic as possible, while the discriminator tries to distinguish the images output by the generator from the real training images.

In the case of our model, the goal of the generator is, therefore, to learn how to generate a realistic flooded image, whereas the discriminator aims to distinguish whether the images it sees are made by the generator or are actual flooded images. This process leads the generator to gradually improve the quality of the generated images, and to be able to fool the discriminator more and more often. At the end of the GAN training process, the images created by the

generator should be indistinguishable from the real images by the discriminator: this is called *convergence*. In practice, convergence is complex to attain, since images have many characteristics and attributes, and so the generator has to learn a very complex set of mathematical functions in order to generate realistic images. This is why we had to try several different approaches in order to be able to generate realistic images of flooding.

WHERE AND HOW TO FLOOD

Initially, we attempted generating flooded scenes using a CycleGAN network,²⁸ a generative model that was shown to be successful in carrying out simple domain transfer manipulations such as transforming horses to zebras and apples to oranges. However, we realized that this relatively simple approach had several drawbacks: on the one hand, it changed the entire image, modifying the appearance of houses, trees, cars, and sky, rather than focusing on the actual flooded region. On the other hand, the way in which CycleGAN functions requires an image to be translated back to its original domain once it has been translated first time, preventing major destructive changes to the initial translation, such as hiding the floor with water: because such a transformation removes information the original image can hardly be reconstructed perfectly, suggesting other approaches were necessary.

We also quickly realized that the task of flooding a scene is complex, since the model being used should learn how to do several things, such as: 1) to put water in sensible areas (and therefore understand where the ground is, the size and perspective of objects, etc.) and 2) to draw realistic water (with reflections of its surroundings, the appropriate texture etc.). We therefore endeavored to create our own, custom AI model. Inspired by previous work such as InstaGAN¹⁹ and Attention-Guided CycleGAN,¹⁷ we divided the flooding task problem into the two subtasks defined above: the *masking* task, which aims to produce masks of regions where water should go, and the *painting* task, which aims to paint realistic water in the regions defined by these masks. However, at this time, we were faced with a problem that comes up often in AI applications: a lack of sufficient data to train our model. We therefore ventured to create a simulated world that we could use to generate more training images for our model.

USING SIMULATED DATA TO LEARN

Generally speaking, AI approaches require a lot of data. This is due to the fact that an AI model needs to see




FIGURE 3. Before and after flooded scene from our Unity simulated world.

many examples of a given class to learn what defines it. This applies to many applications of AI, be it classifying cats and dogs of different breeds or recognizing spoken words with different accents. Generative models can require even more data than traditional models, depending on the complexity of the generative task: sometimes tens or even hundreds of thousands of images are used for training GANs that can generate realistic, high-definition results. While it is relatively easy to gather many examples of faces, zebras, or apples, in our use case, real images of flooding taken from a first-person perspective (as opposed to aerial photographs) are rare, especially in the case of extreme flooding. Also, the images that we were able to painstakingly collect by hand and via crowd-sourcing were heterogeneous in terms of quality and lacking in potentially useful information such as the quantity of water present and the dimensions of objects in the image.

Therefore, in order to address the lack of before/after flooding data, we created a simulated world using the Unity 3D game engine (Unity 3D, engine version 2018.2.21f1) that we could use to generate realistic images. By creating and extending this virtual world, we could generate a wide variety of images for training our model, including urban and suburban areas with different types of buildings and landscapes, covering a variety of homes and buildings that exist worldwide. Using a simulated world also enables us to

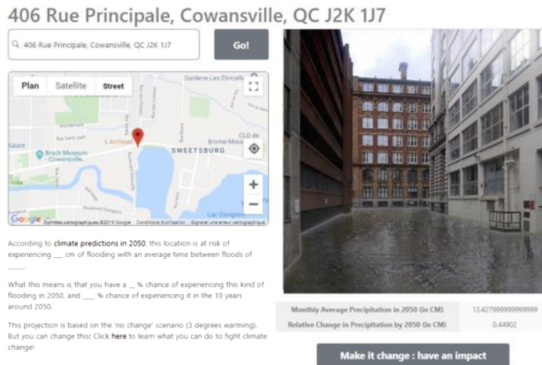


FIGURE 4. Screenshot of the website that will host our AI climate impact visualizer.

access information that is missing in real-world data, for instance pairs of images (i.e., the exact same location before and after a flood, see Fig. 3) and 3-D geometry and depth information of the scene, which is important to enable our model to learn what constitutes different levels of flooding.

Based on a larger dataset including both real and simulated data, we are able to train our GAN architecture, which is composed of the two parts that we mentioned above: a *Masker* model whose task is to generate a mask of where flooding should be on a given image leveraging our simulated data using a technique called ADVENT,²⁶ and a *Painter* model whose aim is to generate realistic water on the area defined by the Masker, inspired by the GauGAN network.²⁰ The *painter* is trained using real images of floods, in which it should be capable of reconstructing water that is masked out using the masker. The model will therefore learn not only what constitutes realistic water, but also the contextual nature of the phenomenon, including shadows and reflections. The parts of the image that are not affected by flooding (i.e., the buildings and the sky) are reconstructed by the network in a manner that is as similar as possible to the input image.

AI CLIMATE IMPACT VISUALIZER

While the AI component is a central part of our endeavor, the ultimate goal of our project is to create a website allowing a user to enter an address anywhere in the world, be it their house, their school, or their workplace, and provide them with an AI-generated imagined rendering of a possible future of this location in 2050 following the detrimental effects of climate change. This image will be accompanied by information regarding the science behind climate change on why extreme weather events are becoming

more frequent and what kinds of changes are happening on a local and global scale. While it is impossible to predict whether a particular location will actually be hit by a climate change-related extreme event, such a visualization aims at encouraging empathy and making users better understand, even if they themselves are not vulnerable to the impacts of climate change, what other people may go through due to extreme weather events. We are also currently working on supplementary AI models to represent smog, wildfires, and droughts in addition to floods, to represent the diversity of extreme weather events that will be amplified by climate change.²⁴

The computer infrastructure behind this website will enable the following to happen: based on an address entered by the user, the system will query Google Streetview to fetch a first-person image of the address, then send this image to the Masker, which will create a mask of where the flood should be drawn. The mask will then be provided to the Painter model along with the original image, outputting a visualization of what a flood could potentially look like at that particular location, the whole generative process taking no more than a few seconds. We will accompany the image with accessible explanations of the science behind climate change, as well as examples of personal and collective actions that can be taken to reduce our environmental footprint. Our aim with this is to empower users of our website, showing them that it is not too late to change the course of climate change and that there are actions that can be taken to fight climate change, and that both as individuals and as members of society, we can do our part to mitigating the extreme climate events that may take place in the future.

CONCLUSION

Climate change is a cognitively overwhelming phenomenon that makes it difficult for us to take action on. In fact, we are hardwired by our genes and by our surroundings to maintain business as usual, and to address threats that are both more material and more urgent than this abstract, diaphanous entity. Using recent progress in artificial intelligence, we are making the impacts of climate change more concrete and more personal, with the goal of helping individuals take action and to make an impact.

We are currently pursuing work on two fronts: we are carrying out lab-based experimentation to establish whether the images we generate can be used to influence the perception of climate risks and public opinion,¹⁸ and we are working on creating a website to

host our AI models and to create a seamless user experience for our audience, including a narrative around the future risks of climate change and extreme weather events, in order to better explain the science behind the predictions and to better contextualize our tool. Our aim is that by early 2021, we will create an immersive, impactful website that will enable our viewers to render possible futures more concrete, conceptualize the impacts of climate change, and make a difference in stopping it in its tracks.

REFERENCES

1. S. J. Ahn, "Embodied experiences in immersive virtual environments: Effects on pro-environmental attitude and behavior," Doctoral dissertation, Stanford Univ., Stanford, CA, USA, 2011.
2. J. Angel, A. LaValle, D. M. Iype, S. Sheppard, and A. Dulic, "Future Delta 2.0 an experiential learning context for a serious game about local climate change," in *Proc. SIGGRAPH Asia Symp. Educ.*, 2015, Art. no. 12.
3. K. Bousmalis, N. Silberman, D. Dohan, D. Erhan, and D. Krishnan, "Unsupervised pixel-level domain adaptation with generative adversarial networks," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2017, pp. 3722–3731.
4. D. A. Chapman, A. Corner, R. Webster, and E. M. Markowitz, "Climate visuals: A mixed methods investigation of public perceptions of climate images in three countries," *Global Environ. Change*, vol. 41, pp. 172–182, Nov. 2016.
5. L. C. Chen, G. Papandreou, I. Kokkinos, K. Murphy, and A. L. Yuille, "DeepLab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected CRFs," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 40, no. 4, pp. 834–848, Apr. 2018.
6. A. Corner and J. Clarke, *Talking Climate: From Research to Practice in Public Engagement*. New York, NY, USA: Springer, 2016.
7. D. Eigen, C. Puhrsch, and R. Fergus, "Depth map prediction from a single image using a multi-scale deep network," in *Proc. Adv. Neural Inf. Process. Syst.*, 2014, pp. 2366–2374.
8. M. C. Fitzpatrick and R. R. Dunn, "Contemporary climatic analogs for 540 North American urban areas in the late 21st century," *Nature Commun.*, vol. 10, no. 1, 2019, Art. no. 614.
9. G. Giannachi, "Representing, performing and mitigating climate change in contemporary art practice," *Leonardo*, vol. 45, no. 2, pp. 124–131, 2012.
10. D. G. Goldstein, H. E. Hershfield, and S. Benartzi, "The illusion of wealth and its reversal," *J. Market. Res.*, vol. 53, no. 5, pp. 804–813, 2016.
11. I. Goodfellow et al., "Generative adversarial nets," in *Proc. Adv. Neural Inf. Process. Syst.*, 2014, pp. 2672–2680.
12. H. E. Hershfield et al., "Increasing saving behavior through age-progressed renderings of the future self," *J. Market. Res.*, vol. 48, pp. S23–S37, 2011.
13. X. Huang, M. Y. Liu, S. Belongie, and J. Kautz, "Multimodal unsupervised image-to-image translation," in *Proc. Eur. Conf. Comput. Vision*, 2018, pp. 172–189.
14. P. Isola, J. Y. Zhu, T. Zhou, and A. A. Efros, "Image-to-image translation with conditional adversarial networks," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2017, pp. 1125–1134.
15. D. Johnson and S. Levin, "The tragedy of cognition: Psychological biases and environmental inaction," *Current Sci.*, vol. 97, no. 11, pp. 1593–1603, 2009. Accessed: Jun. 23, 2020 [Online]. Available: www.jstor.org/stable/24107300
16. A. Luccioni and Y. Bengio, "On the morality of artificial intelligence [commentary]," *IEEE Technol. Soc. Mag.*, vol. 39, no. 1, pp. 16–25, Mar. 2020, doi: [10.1109/MTS.2020.2967486](https://doi.org/10.1109/MTS.2020.2967486).
17. Y. A. Mejjati, C. Richardt, J. Tompkin, D. Cosker, and K. I. Kim, "Unsupervised attention-guided image-to-image translation," in *Proc. Adv. Neural Inf. Process. Syst.*, 2018, pp. 3693–3703.
18. M. Mildenerger, P. Howe, E. Lachapelle, L. Stokes, J. Marlon, and T. Gravelle, "The distribution of climate change public opinion in Canada," *PLoS one*, vol. 11, no. 8, 2016, Art. no. e0159774.
19. S. Mo, M. Cho, and J. Shin, "InstaGAN: Instance-aware image-to-image translation," in *Proc. Int. Conf. Learn. Representations*, 2018.
20. T. Park, M. Y. Liu, T. C. Wang, and J. Y. Zhu, "Semantic image synthesis with spatially-adaptive normalization," in *Proc. IEEE Conf. Comput. Vision Pattern Recognit.*, 2019, pp. 2337–2346.
21. K. Riahi et al., "RCP 8.5—A scenario of comparatively high greenhouse gas emissions," *Climatic Change*, vol. 109, no. 1/2, 2011, Art. no. 33.
22. D. Rolnick et al., "Tackling climate change with machine learning," 2019, [arXiv:1906.05433](https://arxiv.org/abs/1906.05433).
23. S. R. Sheppard, *Visualizing Climate Change: A Guide to Visual Communication of Climate Change and Developing Local Solutions*. Abingdon, U.K.: Routledge, 2012.
24. P. Stott, "Weather risks in a warming world," *Nature Climate Change*, vol. 5, no. 6, pp. 516–517, 2015.

25. S. Van der Linden, E. Maibach, and A. Leiserowitz, "Improving public engagement with climate change: five 'best practice' insights from psychological science," *Perspect. Psychol. Sci.*, vol. 10, pp. 758–763, 2015.
26. T. H. Vu, H. Jain, M. Bucher, M. Cord, and P. Pérez, "Advent: Adversarial entropy minimization for domain adaptation in semantic segmentation," in *Proc. IEEE Conf. Comput. Vision Pattern Recognit.*, 2019, pp. 2517–2526.
27. S. Wang, A. Corner, D. Chapman, and E. Markowitz, "Public engagement with climate imagery in a changing digital landscape," *Wiley Interdisciplinary Rev. Climate Change*, vol. 9, no. 2, 2018, Art. no. e509.
28. J. Y. Zhu, T. Park, P. Isola, and A. A. Efros, "Unpaired image-to-image translation using cycle-consistent adversarial networks," in *Proc. IEEE Int. Conf. Comput. Vision*, 2017, pp. 2223–2232.

ALEXANDRA LUCCIONI is a Postdoctoral Researcher with Mila, Université de Montréal, QC, Canada, under the supervision of Yoshua Bengio. She works on projects that aim to use Artificial Intelligence in "for good" applications such as climate change, health, ethics, and education. She is the corresponding author of this article. Contact her at sasha.luccioni@mila.quebec.

VICTOR SCHMIDT is working toward the Ph.D. degree with Mila, Université de Montréal, QC, Canada, under the supervision of Yoshua Bengio. His research focuses on how AI can contribute to climate change mitigation. Contact him at schmidt@mila.quebec.

VAHE VARDANYAN is a Research & Development Scientist working on computer vision related projects at Ubisoft La Forge, Montreal, QC, Canada. Partially, he is collaborating with Mila on the "Climate Change AI" Project, creating datasets of simulated images using game engines. Contact him at vardanyan.vahe@gmail.com.

YOSHUA BENGIO is a Full Professor with the Department of Computer Science and Operations Research and the Canada Research Chair in Statistical Learning Algorithms, Université de Montréal, as well as the Scientific Director of the Mila Institute. He is a renowned expert in deep learning with a research focus on the principles of learning that yield intelligence. He is a recipient of the 2018 Turing Award, known as the "Nobel Prize of computing," for his work on deep neural networks. Contact him at Yoshua.Bengio@mila.quebec.

Contact department editor Theresa-Marie Rhyne at theresa-marierhyne@gmail.com.